MIST: l_0 Sparse Linear Regression with Momentum

Goran Marjanovic, Magnus O. Ulfarsson, † Alfred O. Hero III ‡

arXiv:1409.7193v2 [stat.ML] 19 Mar 2015

***Abstract*—Significant attention has been given to minimizing a penalized least squares criterion for estimating sparse solutions to large linear systems of equations. The penalty is responsible for inducing sparsity and the natural choice is the so-called l_0 norm. In this paper we develop a Momentumized Iterative Shrinkage Thresholding (MIST) algorithm for minimizing the resulting non-convex criterion and prove its convergence to a local minimizer. Simulations on large data sets show superior performance of the proposed method to other methods.**

***Index Terms*—sparsity, non-convex, l_0 regularization, linear regression, iterative shrinkage thresholding, hard-thresholding, momentum**

I. Introduction

In the current age of big data acquisition there has been an ever growing interest in sparse representations, which consists of representing, say, a noisy signal as a linear combination of very few components. This implies that the entire information in the signal can be approximately captured by a small number of components, which has huge benefits in analysis, processing and storage of high dimensional signals. As a result, sparse linear regression has been widely studied with many applications in signal and image processing, statistical inference and machine learning. Specific applications include compressed sensing, denoising, inpainting, deblurring, source separation, sparse image reconstruction, and signal classification, etc.

The linear regression model is given by:

$$\mathbf{y} = \mathbf{A}\mathbf{x} + \boldsymbol{\epsilon}, \tag{1}$$

where $\mathbf{y}_{d\times 1}$ is a vector of noisy data observations, $\mathbf{x}_{m\times 1}$ is the sparse representation (vector) of interest, $\mathbf{A}_{d\times m}$ is the regression matrix and $\boldsymbol{\epsilon}_{d\times 1}$ is the noise. The estimation aim is to choose the simplest model, i.e., the sparsest $\mathbf{x}$, that adequately explains the data $\mathbf{y}$. To estimate a sparse $\mathbf{x}$, major attention has been given to minimizing a sparsity Penalized Least Squares (PLS) criterion [1]–[10]. The least squares term promotes goodness-of-fit of the estimator while the penalty shrinks its coefficients to zero. Here we consider the non-convex l_0 penalty since it is the natural sparsity promoting penalty and induces maximum sparsity. The resulting non-convex l_0 PLS criterion is given by:

$$F(\mathbf{x}) = \frac{1}{2}\|\mathbf{y} - \mathbf{A}\mathbf{x}\|_2^2 + \lambda\|\mathbf{x}\|_0, \tag{2}$$

where $\lambda > 0$ is the tuning/regularization parameter and $\|\mathbf{x}\|_0$ is the l_0 penalty representing the number of non-zeros in $\mathbf{x}$.

A. Previous Work

Existing algorithms for directly minimizing (2) fall into the category of Iterative Shrinkage Thresholding (IST), and rely on the Majorization-Minimization (MM) type procedures, see [1], [2]. These procedures exploit separability properties of the l_0 PLS criterion, and thus, rely on the minimizers of one dimensional versions of the PLS function: the so-called hard-thresholding operators. Since the convex l_1 PLS criterion has similar separability properties, some MM procedures developed for its minimization could with modifications be applied to minimize (2). Applicable MM procedures include first order methods and their accelerated versions [9], [11], [12]. However, when these are applied to the l_0 penalized problem (2) there is no guarantee of convergence, and for [9] there is additionally no guarantee of algorithm stability.

Analysis of convergence of MM algorithms for minimizing the l_0 PLS criterion (2) is rendered difficult due to lack of convexity. As far as we are aware, algorithm convergence for this problem has only been shown for the Iterative Hard Thresholding (IHT) method [1], [2]. Specifically, a bounded sequence generated by IHT was shown to converge to the set of local minimizers of (2) when the singular values of $\mathbf{A}$ are strictly less than one. Convergence analysis of algorithms designed for minimizing the l_q PLS criterion, $q \in (0, 1]$, is not applicable to the case of the l_0 penalized objective (2) because it relies on convex arguments when $q = 1$, and continuity and/or differentiability of the criterion when $q \in (0, 1)$.

B. Paper Contribution

In this paper we develop an MM algorithm with momentum acceleration, called Momentumized IST (MIST), for minimizing the l_0 PLS criterion (2) and prove its convergence to a single local minimizer without imposing any assumptions on $\mathbf{A}$. Simulations on large data sets are carried out, which show that the proposed algorithm outperforms existing methods for minimizing (2), including modified MM methods originally designed for the l_1 PLS criterion.

The paper is organised as follows. Section II reviews some of background on MM that will be used to develop the proposed convergent algorithm. The proposed algorithm is given in Section III, and Section IV contains the convergence

Goran Marjanovic is with the School of Electrical Engineering and Telecommunications, University of New South Wales, Sydney, NSW, 2052 Australia (email: goran.m.marjanovic@gmail.com)

Alfred O. Hero III is with the Department of Electrical and Computer Science, University of Michigan, Ann Arbor, MI, 48109 USA (email: hero@eecs.umich.edu)

Magnus O. Ulfarsson is with the Department of Electrical Engineering, University of Iceland, Reykjavik, 107 Iceland (email: mou@hi.is)

† This work was partly supported by the Research Fund of the University of Iceland and the Icelandic Research Fund (130635-051).

‡ This work was partially supported by ARO grant W911NF-11-1-0391.

analysis. Lastly, Section V and VI presents the simulations and concluding remarks respectively.

C. Notation

The i-th component of a vector $\mathbf{v}$ is denoted by $\mathbf{v}[i]$. Given a vector $f(\mathbf{v})$, where $f(\cdot)$ is a function, its i-th component is denoted by $f(\mathbf{v})[i]$. $\|\mathbf{M}\|$ is the spectral norm of matrix $\mathbf{M}$. $\mathbb{I}(\cdot)$ is the indicator function equaling 1 if its argument is true, and zero otherwise. Given a vector $\mathbf{v}$, $\|\mathbf{v}\|_0 = \sum_i \mathbb{I}(\mathbf{v}[i] \neq 0)$. sgn$(\cdot)$ is the sign function. $\{\mathbf{x}_k\}_{k\geq 0}$ denotes an infinite sequence, and $\{\mathbf{x}_{k_n}\}_{n\geq 0}$ an infinite subsequence, where $k_n \leq k_{n+1}$ for all $n \geq 0$.

II. Preliminaries

Denoting the least squares term in (2) by:

$$f(\mathbf{x}) = \frac{1}{2}\|\mathbf{y} - \mathbf{A}\mathbf{x}\|_2^2,$$

the Lipschitz continuity of $\nabla f(\cdot)$ implies:

$$f(\mathbf{z}) \leq f(\mathbf{x}) + \nabla f(\mathbf{x})^T(\mathbf{z} - \mathbf{x}) + \frac{\mu}{2}\|\mathbf{z} - \mathbf{x}\|_2^2$$

for all $\mathbf{x}, \mathbf{z}, \mu \geq \|\mathbf{A}\|^2$. For the proof see [9, Lemma 2.1]. As a result, the following approximation of the objective function $F(\cdot)$ in (2),

$$Q_\mu(\mathbf{z}, \mathbf{x}) = f(\mathbf{x}) + \nabla f(\mathbf{x})^T(\mathbf{z} - \mathbf{x}) + \frac{\mu}{2}\|\mathbf{z} - \mathbf{x}\|_2^2 + \lambda\|\mathbf{z}\|_0 \quad (3)$$

is a majorizing function, i.e.,

$$F(\mathbf{z}) \leq Q_\mu(\mathbf{z}, \mathbf{x}) \text{ for any } \mathbf{x}, \mathbf{z}, \mu \geq \|\mathbf{A}\|^2. \quad (4)$$

Let $\mathcal{P}_\mu(\mathbf{x})$ be any point in the set $\arg\min_{\mathbf{z}} Q_\mu(\mathbf{z}, \mathbf{x})$, we have:

$$F(\mathcal{P}_\mu(\mathbf{x})) \overset{(4)}{\leq} Q_\mu(\mathcal{P}_\mu(\mathbf{x}), \mathbf{x}) \leq Q_\mu(\mathbf{x}, \mathbf{x}) = F(\mathbf{x}), \quad (5)$$

where the stacking of (4) above the first inequality indicates that this inequality follows from Eq. (4). The proposed algorithm will be constructed based on the above MM framework with a momentum acceleration, described below. This momentum acceleration will be designed based on the following.

Theorem 1. Let $\mathbf{B}_\mu = \mu\mathbf{I} - \mathbf{A}^T\mathbf{A}$, where $\mu > \|\mathbf{A}\|^2$, and:

$$\alpha = 2\eta\left(\frac{\boldsymbol{\delta}^T\mathbf{B}_\mu(\mathcal{P}_\mu(\mathbf{x}) - \mathbf{x})}{\boldsymbol{\delta}^T\mathbf{B}_\mu\boldsymbol{\delta}}\right),\ \eta \in [0, 1], \quad (6)$$

where $\boldsymbol{\delta} \neq 0$. Then, $F\left(\mathcal{P}_\mu(\mathbf{x} + \alpha\boldsymbol{\delta})\right) \leq F(\mathbf{x})$.

For the proof see the Appendix.

A. Evaluating the Operator $\mathcal{P}_\mu(\cdot)$

Since (3) is non-convex there may exist multiple minimizers of $Q_\mu(\mathbf{z}, \cdot)$ so that $\mathcal{P}_\mu(\cdot)$ may not be unique. We select a single element of the set of minimizers as described below. By simple algebraic manipulations of the quadratic quantity in (3), letting:

$$g(\mathbf{x}) = \mathbf{x} - \frac{1}{\mu}\nabla f(\mathbf{x}), \quad (7)$$

it is easy to show that:

$$Q_\mu(\mathbf{z}, \mathbf{x}) = \frac{\mu}{2}\|\mathbf{z} - g(\mathbf{x})\|_2^2 + \lambda\|\mathbf{z}\|_0 + f(\mathbf{x}) - \frac{1}{2\mu}\|\nabla f(\mathbf{x})\|_2^2,$$

and so, $\mathcal{P}_\mu(\cdot)$ is given by:

$$\mathcal{P}_\mu(\mathbf{x}) = \arg\min_{\mathbf{z}}\ \frac{1}{2}\|\mathbf{z} - g(\mathbf{x})\|_2^2 + (\lambda/\mu)\|\mathbf{z}\|_0. \quad (8)$$

For the proposed algorithm we fix $\mathcal{P}_\mu(\cdot) = \mathcal{H}_{\lambda/\mu}(g(\cdot))$, the point to point map defined in the following Theorem.

Theorem 2. Let the hard-thresholding (point-to-point) map $\mathcal{H}_h(\cdot)$, $h > 0$, be such that for each $i = 1, \ldots, m$:

$$\mathcal{H}_h(g(\mathbf{v}))[i] = \begin{cases} 0 & \text{if } |g(\mathbf{v})[i]| < \sqrt{2h} \\ g(\mathbf{v})[i]\ \mathbb{I}(\mathbf{v}[i] \neq 0) & \text{if } |g(\mathbf{v})[i]| = \sqrt{2h} \\ g(\mathbf{v})[i] & \text{if } |g(\mathbf{v})[i]| > \sqrt{2h}. \end{cases} \quad (9)$$

Then, $\mathcal{H}_{\lambda/\mu}(g(\mathbf{v})) \in \arg\min_{\mathbf{z}} Q_\mu(\mathbf{z}, \mathbf{v})$, where $g(\cdot)$ is from (7).

The proof is in the Appendix.

Evidently Theorem 1 holds with $\mathcal{P}_\mu(\cdot)$ replaced by $\mathcal{H}_{\lambda/\mu}(g(\cdot))$. The motivation for selecting this particular minimizer is Lemma 2 in Section IV.

III. The Algorithm

The proposed algorithm is constructed by repeated application of Theorem 1 where $\boldsymbol{\delta}$ is chosen to be the difference between the current and the previous iterate, i.e.,

$$\mathbf{x}_{k+1} = \mathcal{H}_{\lambda/\mu}\left(\mathbf{w}_k - \frac{1}{\mu}\nabla f(\mathbf{w}_k)\right),\ \mathbf{w}_k = \mathbf{x}_k + \alpha_k\boldsymbol{\delta}_k \quad (10)$$

with α_k given by (6), where $\boldsymbol{\delta}_k = \mathbf{x}_k - \mathbf{x}_{k-1}$. The iteration (10) is an instance of a momentum accelerated IST algorithm, similar to Fast IST Algorithm (FISTA) introduced in [9] for minimizing the convex l_1 PLS criterion. In (10), $\boldsymbol{\delta}_k$ is called the momentum term and α_k is a momentum step size parameter. A more explicit implementation of (10) is given below. Our proposed algorithm will be called Momentumized Iterative Shrinkage Thresholding (MIST).

Momentumized IST (MIST) Algorithm

Compute $\bar{\mathbf{y}} = (\mathbf{y}^T\mathbf{A})^T$ off-line. Choose $\mathbf{x}_0$ and let $\mathbf{x}_{-1} = \mathbf{x}_0$. Also, calculate $\|\mathbf{A}\|^2$ off-line, let $\mu > \|\mathbf{A}\|^2$ and $k = 0$. Then:

(1) If $k = 0$, let $\alpha_k = 0$. Otherwise, compute:
  (a) $\mathbf{u}_k = \mathbf{A}\mathbf{x}_k$
  (b) $\mathbf{v}_k = (\mathbf{u}_k^T\mathbf{A})^T$
  (c) $\mathbf{g}_k = \mathbf{x}_k - \frac{1}{\mu}(\mathbf{v}_k - \bar{\mathbf{y}})$
  (d) $\mathbf{p}_k = \mathcal{H}_{\lambda/\mu}(\mathbf{g}_k) - \mathbf{x}_k$
  (e) $\boldsymbol{\delta}_k = \mathbf{x}_k - \mathbf{x}_{k-1}$ and:

$$\boldsymbol{\gamma}_k = \mu\boldsymbol{\delta}_k - \mathbf{v}_k + \mathbf{v}_{k-1} \quad (11)$$

  (f) Choose $\eta_k \in (0, 1)$ and compute:

$$\alpha_k = 2\eta_k\left(\frac{\boldsymbol{\gamma}_k^T\mathbf{p}_k}{\boldsymbol{\gamma}_k^T\boldsymbol{\delta}_k}\right) \quad (12)$$

(2) Using (c), (e) and (f) compute:

$$\mathbf{x}_{k+1} = \mathcal{H}_{\lambda/\mu}\left(\mathbf{g}_k + \frac{\alpha_k}{\mu}\boldsymbol{\gamma}_k\right) \tag{13}$$

(3) Let $k = k + 1$ and go to (1).

Remark 1. Thresholding using (9) is simple, and can always be done off-line. Secondly, note that MIST requires computing only $\mathcal{O}(2md)$ products, which is the same order required when the momentum term $\boldsymbol{\delta}_k$ is not incorporated, i.e., $\eta_k = 0$ for all k. In this case, MIST is a generalization of IHT from [1], [2]. Other momentum methods such as FISTA [9] and its monotone version M-FISTA [11] also require computing $\mathcal{O}(2md)$ and $\mathcal{O}(3md)$ products, respectively.

IV. Convergence Analysis

Here we prove that the proposed MIST algorithm iterates converge to a local minimizer of $F(\cdot)$.

Theorem 3. Suppose $\{\mathbf{x}_k\}_{k\geq 0}$ is a bounded sequence generated by the MIST algorithm. Then $\mathbf{x}_k \to \mathbf{x}_\bullet$ as $k \to \infty$, where $\mathbf{x}_\bullet$ is a local minimizer of (2).

The proof is in the Appendix and requires several lemmas that are also proved in the Appendix.

In Lemma 1 and 2 it is assumed that MIST reaches a fixed point only in the limit, i.e., $\mathbf{x}_{k+1} \neq \mathbf{x}_k$ for all k. This implies that $\boldsymbol{\delta}_k \neq 0$ for all k.

Lemma 1. $\mathbf{x}_{k+1} - \mathbf{x}_k \to 0$ as $k \to \infty$.

The following lemma motivates Theorem 2 and is crucial for the subsequent convergence analysis.

Lemma 2. Assume the result in Lemma 1. If, for any subsequence $\{\mathbf{x}_{k_n}\}_{n\geq 0}$, $\mathbf{x}_{k_n} \to \mathbf{x}_\bullet$ as $n \to \infty$, then:

$$\mathcal{H}_{\lambda/\mu}\left(\mathbf{w}_{k_n} - \frac{1}{\mu}\nabla f(\mathbf{w}_{k_n})\right) \to \mathcal{H}_{\lambda/\mu}\left(\mathbf{x}_\bullet - \frac{1}{\mu}\nabla f(\mathbf{x}_\bullet)\right), \tag{14}$$

where $\mathbf{w}_{k_n} = \mathbf{x}_{k_n} + \alpha_{k_n}\boldsymbol{\delta}_{k_n}$.

The following lemma characterizes the fixed points of the MIST algorithm.

Lemma 3. Suppose $\mathbf{x}_\bullet$ is a fixed point of MIST. Letting $\mathcal{Z} = \{i : \mathbf{x}_\bullet[i] = 0\}$ and $\mathcal{Z}^c = \{i : \mathbf{x}_\bullet[i] \neq 0\}$,

(C_1) If $i \in \mathcal{Z}$, then $|\nabla f(\mathbf{x}_\bullet)[i]| \leq \sqrt{2\lambda\mu}$.

(C_2) If $i \in \mathcal{Z}^c$, then $\nabla f(\mathbf{x}_\bullet)[i] = 0$.

(C_3) If $i \in \mathcal{Z}^c$, then $|\mathbf{x}_\bullet[i]| \geq \sqrt{2\lambda/\mu}$.

Lemma 4. Suppose $\mathbf{x}_\bullet$ is a fixed point of MIST. Then there exists $\epsilon > 0$ such that $F(\mathbf{x}_\bullet) < F(\mathbf{x}_\bullet + \mathbf{d})$ for any $\mathbf{d}$ satisfying $\|\mathbf{d}\|_2 \in (0, \epsilon)$. In other words, $\mathbf{x}_\bullet$ is a strict local minimizer of (2).

Lemma 5. The limit points of $\{\mathbf{x}_k\}_{k\geq 0}$ are fixed points of MIST.

All of the above lemmas are proved in the Appendix.

V. Simulations

Here we demonstrate the performance advantages of the proposed MIST algorithm in terms of convergence speed. The methods used for comparison are the well known MM algorithms: ISTA and FISTA from [9], as well as M-FISTA from [11], where the soft-thresholding map is replaced by the hard-thresholding map. In this case, ISTA becomes identical to the IHT algorithm from [1], [2], while FISTA and M-FISTA become its accelerated versions, which exploit the ideas in [13].

A popular compressed sensing scenario is considered with the aim of reconstructing a length m sparse signal $\mathbf{x}$ from d observations, where $d < m$. The matrix $\mathbf{A}_{d\times m}$ is obtained by filling it with independent samples from the standard Gaussian distribution. A relatively high dimensional example is considered, where $d = 2^{13} = 8192$ and $m = 2^{14} = 16384$, and $\mathbf{x}$ contains 150 randomly placed ± 1 spikes (0.9% non-zeros). The observation $\mathbf{y}$ is generated according to (1) with the standard deviation of the noise $\boldsymbol{\epsilon}$ given by $\sigma = 3, 6, 10$. The Signal to Noise Ratio (SNR) is defined by:

$$\text{SNR} = 10\log_{10}\left(\frac{\|\mathbf{A}\mathbf{x}\|_2^2}{\sigma^2 d}\right).$$

Figures 1, 2 and 3 show a plot of $\mathbf{A}\mathbf{x}$ and observation noise $\boldsymbol{\epsilon}$ for the three SNR values corresponding to the three considered values of σ.

A. Selection of the Tuning Parameter λ

Oracle results are reported, where the chosen tuning parameter λ in (2) is the minimizer of the Mean Squared Error (MSE), defined by:

$$\text{MSE}(\lambda) = \frac{\|\mathbf{x} - \widehat{\mathbf{x}}\|_2^2}{\|\mathbf{x}\|_2^2},$$

where $\widehat{\mathbf{x}} = \widehat{\mathbf{x}}(\lambda)$ is the estimator of $\mathbf{x}$ produced by a particular algorithm.

As $\mathbf{x}$ is generally unknown to the experimenter we also report results of using a model selection method to select λ. Some of the classical model selection methods include the Bayesian Information Criterion (BIC) [14], the Akaike Information criterion [15], and (generalized) cross validation [16], [17]. However, these methods tend to select a model with many spurious components when m is large and d is comparatively smaller, see [18]–[21]. As a result, we use the Extended BIC (EBIC) model selection method proposed in [18], which incurs a small loss in the positive selection rate while tightly controlling the false discovery rate, a desirable property in many applications. The EBIC is defined by:

$$\text{EBIC}(\lambda) = \log\left(\frac{\|\mathbf{y} - \mathbf{A}\widehat{\mathbf{x}}\|_2^2}{d}\right) + \left(\frac{\log d}{d} + 2\gamma\frac{\log m}{d}\right)\|\widehat{\mathbf{x}}\|_0,$$

and the chosen λ in (2) is the minimizer of this criterion. Note that the classical BIC criterion is obtained by setting $\gamma = 0$. As suggested in [18], we let $\gamma = 1 - 1/(2\kappa)$, where κ is the solution of $m = d^\kappa$, i.e.,

$$\kappa = \frac{\log m}{\log d} = 1.08$$

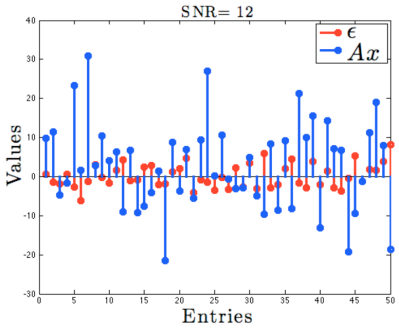

Fig. 1. Examples of the first 50 entry values in the (noiseless) observation $\mathbf{Ax}$ and the observation noise $\boldsymbol{\epsilon}$ for SNR= 12.

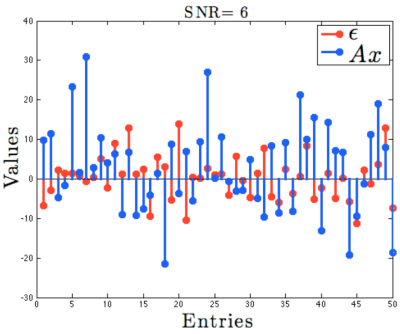

Fig. 2. Examples of the first 50 entry values in the (noiseless) observation $\mathbf{Ax}$ and the observation noise $\boldsymbol{\epsilon}$ for SNR= 6.

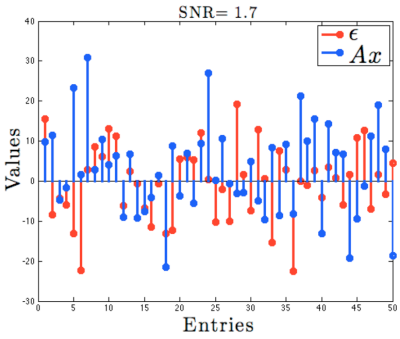

Fig. 3. Examples of the first 50 entry values in the (noiseless) observation $\mathbf{Ax}$ and the observation noise $\boldsymbol{\epsilon}$ for SNR= 1.7.

B. Results

All algorithms are initialized with $\mathbf{x}_0 = \mathbf{0}$, and are terminated when the following criterion is satisfied:

$$\frac{|F(\mathbf{x}_k) - F(\mathbf{x}_{k-1})|}{F(\mathbf{x}_k)} < 10^{-10}. \tag{15}$$

In the MIST algorithm we let $\mu = \|\mathbf{A}\|^2 + 10^{-15}$ and $\eta_k = 1 - 10^{-15}$. All experiments were run in MATLAB 8.1 on an Intel Core i7 processor with 3.0GHz CPU and 8GB of RAM.

Figures 4, 5 and 6 show percentage reduction of $F(\cdot)$ as a function of time and iteration for each algorithm. To make the comparisons fair, i.e., to make sure all the algorithms minimize the same objective function, a common λ is used and chosen to be the smallest λ from the averaged $\arg\min_\lambda \text{EBIC}(\lambda)$ obtained by each algorithm (over 10 instances).

Figures 7, 8 and 9 also show percentage reduction of $F(\cdot)$ as a function of time and iteration for each algorithm. This time the MSE is used for the tuning of λ, and again, to make sure all the algorithms minimize the same objective function, a common λ is used and chosen to be the smallest λ from the averaged $\arg\min_\lambda \text{MSE}(\lambda)$ obtained by each algorithm (over 10 instances).

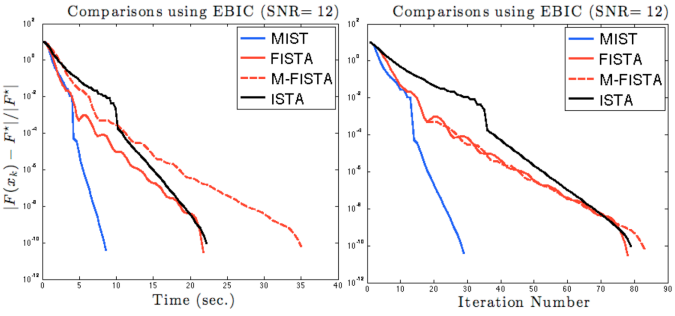

Fig. 4. Algorithm comparisons based on relative error $|F(\mathbf{x}_k) - F^\star|/|F^\star|$ where $F^\star$ is the final value of $F(\cdot)$ obtained by each algorithm at its termination, i.e., $F^\star = F(\mathbf{x}_k)$, where $F(\mathbf{x}_k)$ satisfies the termination criterion in (15). Here SNR=12, and the regularization parameter λ has been selected using the EBIC criterion. As it can be seen, in the low noise environment ($\sigma = 3$) the MIST algorithm outperforms the rest, both in terms of time and iteration number.

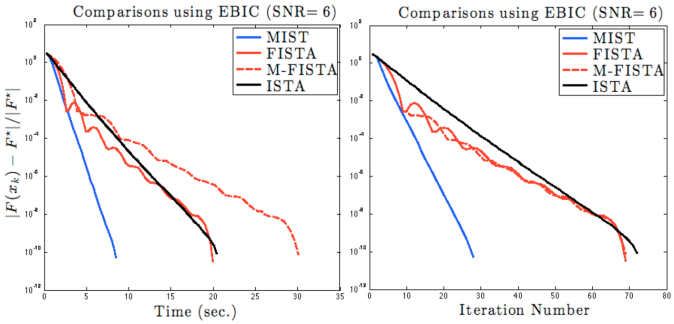

Fig. 5. Similar comparisons as in Fig. 4 except that SNR=6. As it can be seen, in the intermediate noise environment ($\sigma = 6$) the MIST algorithm outperforms the others, both in terms of time and iteration number.

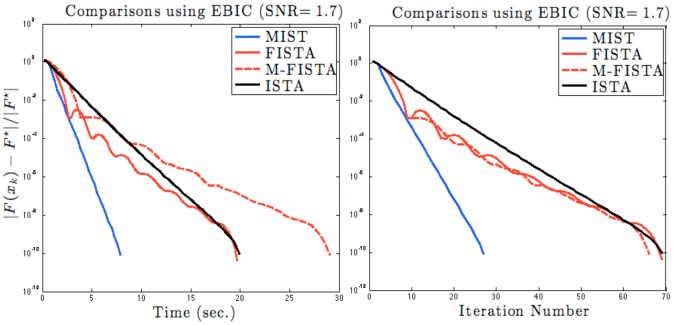

Fig. 6. Similar comparisons as in Fig. 4 except that SNR=1.7. As it can be seen, in the high noise environment ($\sigma = 10$) the MIST algorithm outperforms the rest, both in terms of time and iteration number.

Based on a large number of experiments we noticed that MIST, FISTA and ISTA outperformed M-FISTA in terms of run time. This could be due to the fact that M-FISTA requires computing a larger number of products, see Remark 1, and the fact that it is a monotone version of a severely non-monotone FISTA. The high non-monotonicity could possibly be due to non-convexity of the objective function $F(\cdot)$.

Lastly, Figures 10, 11 and 12 show the average speed of each algorithm as a function of λ.

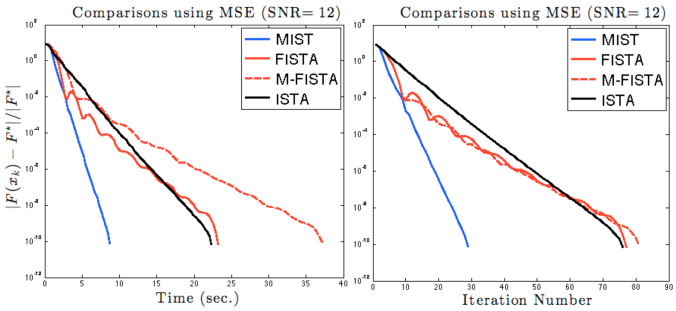

Fig. 7. Algorithm comparisons based on relative error $|F(\mathbf{x}_k) - F^\star|/|F^\star|$ where $F^\star$ is the final value of $F(\cdot)$ obtained by each algorithm at its termination, i.e., $F^\star = F(\mathbf{x}_k)$, where $F(\mathbf{x}_k)$ satisfies the termination criterion in (15). Here an oracle selects the regularization parameter λ using the minimum MSE criterion. As it can be seen, in the low noise environment ($\sigma = 3$) the MIST algorithm outperforms the rest, both in terms of time and iteration number.

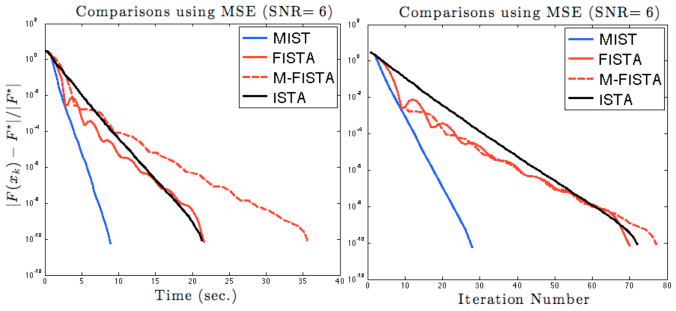

Fig. 8. Similar comparisons as in Fig. 7 except that SNR=6. As it can be seen, in the intermediate noise environment ($\sigma = 6$) the MIST algorithm outperforms the rest, both in terms of time and iteration number.

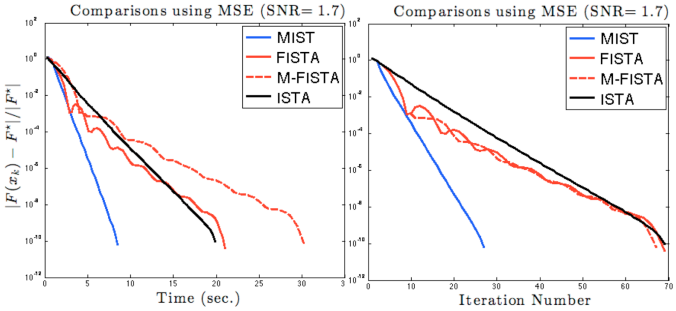

Fig. 9. Similar comparisons as in Fig. 7 except that SNR=1.7. As it can be seen, in the high noise environment ($\sigma = 10$) the MIST algorithm outperforms the rest, both in terms of time and iteration number.

VI. Conclusion

We have developed a momentum accelerated MM algorithm, MIST, for minimizing the l_0 penalized least squares criterion for linear regression problems. We have proved that MIST converges to a local minimizer without imposing any assumptions on the regression matrix $\mathbf{A}$. Simulations on large data sets were carried out for different SNR values, and have shown that the MIST algorithm outperforms other popular MM algorithms in terms of time and iteration number.

VII. Appendix

Proof of Theorem 1: Let $\mathbf{w} = \mathbf{x} + \beta\boldsymbol{\delta}$. The quantities $f(\mathbf{w})$, $\nabla f(\mathbf{w})^T(\mathbf{z} - \mathbf{w})$ and $\|\mathbf{z} - \mathbf{w}\|_2^2$ are quadratic functions, and by simple linear algebra they can easily be expanded in terms of $\mathbf{z}$, $\mathbf{x}$ and $\boldsymbol{\delta}$. Namely,

$$\begin{aligned} f(\mathbf{w}) &= \frac{1}{2}\|\mathbf{A}\mathbf{w} - \mathbf{y}\|^2 \\ &= \frac{1}{2}\|\mathbf{A}\mathbf{x} - \mathbf{y} + \beta\mathbf{A}\boldsymbol{\delta}\|^2 \\ &= f(\mathbf{x}) + \beta(\mathbf{A}\mathbf{x} - \mathbf{y})^T\mathbf{A}\boldsymbol{\delta} + \frac{1}{2}\beta^2\|\mathbf{A}\boldsymbol{\delta}\|^2 \\ &= f(\mathbf{x}) + \beta\nabla f(\mathbf{x})^T\boldsymbol{\delta} + \frac{1}{2}\beta^2\|\mathbf{A}\boldsymbol{\delta}\|^2. \end{aligned} \tag{16}$$

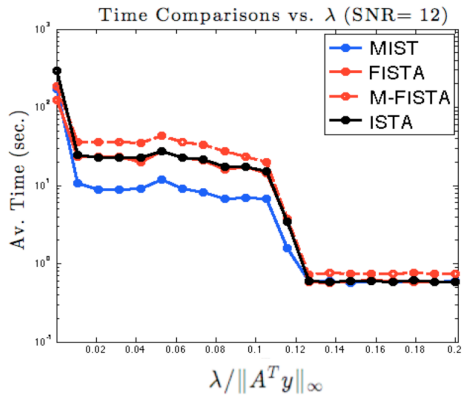

Fig. 10. Average time vs. λ over 5 instances when SNR=12. For the comparison, 20 equally spaced values of λ are considered, where $10^{-4}\|\mathbf{A}^T\mathbf{y}\|_\infty \leq \lambda \leq 0.2\|\mathbf{A}^T\mathbf{y}\|_\infty$. As it can be seen, except in the scenario when $\lambda = 10^{-4}\|\mathbf{A}^T\mathbf{y}\|_\infty$ the MIST algorithm outperforms the others.

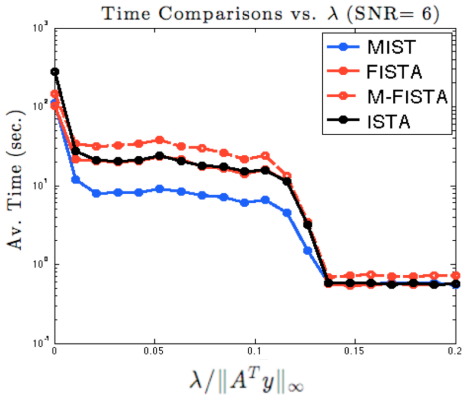

Fig. 11. Similar comparisons as in Fig. 10 except that SNR=6. As it can be seen, except in the scenario when $\lambda = 10^{-4}\|\mathbf{A}^T\mathbf{y}\|_\infty$ the MIST algorithm outperforms the others.

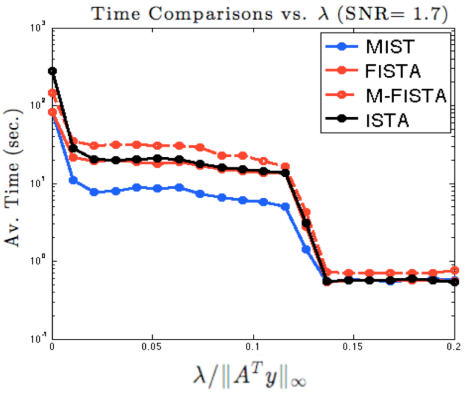

Fig. 12. Similar comparisons as in Fig. 10 except that SNR=1.7. As it can be seen, except in the scenario when $\lambda = 10^{-4}\|\mathbf{A}^T\mathbf{y}\|_\infty$ the MIST algorithm outperforms the others.

Also,

$$\begin{aligned}
&\nabla f(\mathbf{w})^T(\mathbf{z}-\mathbf{w}) \\
&= (\nabla f(\mathbf{x}) + \beta \mathbf{A}^T\mathbf{A}\boldsymbol{\delta})^T(\mathbf{z}-\mathbf{w}) \\
&= \nabla f(\mathbf{x})^T(\mathbf{z}-\mathbf{w}) + \beta\boldsymbol{\delta}^T\mathbf{A}^T\mathbf{A}(\mathbf{z}-\mathbf{w}) \\
&= \nabla f(\mathbf{x})^T(\mathbf{z}-\mathbf{x}-\beta\boldsymbol{\delta}) + \beta\boldsymbol{\delta}^T\mathbf{A}^T\mathbf{A}(\mathbf{z}-\mathbf{x}-\beta\boldsymbol{\delta}) \\
&= \nabla f(\mathbf{x})^T(\mathbf{z}-\mathbf{x}) - \beta\nabla f(\mathbf{x})^T\boldsymbol{\delta} + \beta\boldsymbol{\delta}^T\mathbf{A}^T\mathbf{A}(\mathbf{z}-\mathbf{x}) \\
&\quad - \beta^2\|\mathbf{A}\boldsymbol{\delta}\|^2, \qquad (17)
\end{aligned}$$

and finally:

$$\begin{aligned}
\|\mathbf{z}-\mathbf{w}\|_2^2 &= \|\mathbf{z}-\mathbf{x}-\beta\boldsymbol{\delta}\|_2^2 \\
&= \|\mathbf{z}-\mathbf{x}\|_2^2 - \beta\boldsymbol{\delta}^T(\mathbf{z}-\mathbf{x}) + \beta^2\|\boldsymbol{\delta}\|_2^2 \qquad (18)
\end{aligned}$$

Using the above expansions and the definition of $Q_\mu(\cdot,\cdot)$, we have that:

$$Q_\mu(\mathbf{z},\mathbf{w}) = Q_\mu(\mathbf{z},\mathbf{x}) + \Phi_\mu(\mathbf{z},\boldsymbol{\delta},\beta), \qquad (19)$$

where:

$$\Phi_\mu(\mathbf{z},\boldsymbol{\delta},\beta) = \frac{1}{2}\beta^2\boldsymbol{\delta}^T\mathbf{B}_\mu\boldsymbol{\delta} - \beta\boldsymbol{\delta}^T\mathbf{B}_\mu(\mathbf{z}-\mathbf{x}).$$

Observing that $\boldsymbol{\delta}^T\mathbf{B}_\mu\boldsymbol{\delta} > 0$, let:

$$\beta = 2\eta\left(\frac{\boldsymbol{\delta}^T\mathbf{B}_\mu(\mathbf{z}-\mathbf{x})}{\boldsymbol{\delta}^T\mathbf{B}_\mu\boldsymbol{\delta}}\right),\ \eta\in[0,1]. \qquad (20)$$

Then, one has:

$$\begin{aligned}
Q_\mu(\mathcal{P}_\mu(\mathbf{w}),\mathbf{w}) &= \min_{\mathbf{z}} Q_\mu(\mathbf{z},\mathbf{w}) \le Q_\mu(\mathbf{z},\mathbf{w}) \\
&\overset{(19)}{=} Q_\mu(\mathbf{z},\mathbf{x}) + \Phi_\mu(\mathbf{z},\boldsymbol{\delta},\beta) \\
&\overset{(20)}{=} Q_\mu(\mathbf{z},\mathbf{x}) - 2\eta(1-\eta)\frac{[\boldsymbol{\delta}^T\mathbf{B}_\mu(\mathbf{z}-\mathbf{x})]^2}{\boldsymbol{\delta}^T\mathbf{B}_\mu\boldsymbol{\delta}} \qquad (21) \\
&\le Q_\mu(\mathbf{z},\mathbf{x}), \qquad (22)
\end{aligned}$$

which holds for any $\mathbf{z}$. So, letting $\mathbf{z} = \mathcal{P}_\mu(\mathbf{x})$ implies:

$$F(\mathcal{P}_\mu(\mathbf{w})) \overset{(5)}{\le} Q_\mu(\mathcal{P}_\mu(\mathbf{w}),\mathbf{w}) \overset{(22)}{\le} Q_\mu(\mathcal{P}_\mu(\mathbf{x}),\mathbf{x}) \overset{(5)}{\le} F(\mathbf{x}),$$

which completes the proof. □

Proof of Theorem 2: Looking at (8) it is obvious that:

$$\mathcal{P}_\mu(\mathbf{v})[i] = \arg\min_{\mathbf{z}[i]} \frac{1}{2}(\mathbf{z}[i]-g(\mathbf{v})[i])^2 + (\lambda/\mu)\mathbb{I}(\mathbf{z}[i]\neq 0).$$

If $|g(\mathbf{v})[i]| \neq \sqrt{2\lambda/\mu}$, by [22, Theorem 1] $\mathcal{P}_\mu(\mathbf{v})[i]$ is unique and given by $\mathcal{H}_{\lambda/\mu}(g(\mathbf{v}))[i]$. If $|g(\mathbf{v})[i]| = \sqrt{2\lambda/\mu}$, again by [22, Theorem 1] we now have:

$$\mathcal{P}_\mu(\mathbf{v})[i] = 0 \text{ and } \mathcal{P}_\mu(\mathbf{v})[i] = \text{sgn}(g(\mathbf{v})[i])\sqrt{2\lambda/\mu}.$$

Hence, $\mathcal{H}_{\lambda/\mu}(g(\mathbf{v}))[i] \in \mathcal{P}_\mu(\mathbf{v})[i]$, completing the proof. □

Proof of Lemma 1: From Theorem 1, $0 \le F(\mathbf{x}_{k+1}) \le F(\mathbf{x}_k)$, so the sequence $\{F(\mathbf{x}_k)\}_{k\ge 0}$ is bounded, which means it has a finite limit, say, $F_\bullet$. As a result:

$$F(\mathbf{x}_k) - F(\mathbf{x}_{k+1}) \to F_\bullet - F_\bullet = 0. \qquad (23)$$

Next, recall that $\mathbf{w}_k = \mathbf{x}_k + \alpha_k\boldsymbol{\delta}_k$ and $g(\cdot) = (\cdot) - \frac{1}{\mu}\nabla f(\cdot)$. So, using (21) in the proof of Theorem 1 (where $\mathbf{z} = \mathcal{P}_\mu(\mathbf{x})$) with $\mathbf{w}$, $\mathcal{P}_\mu(\mathbf{w})$, $\mathcal{P}_\mu(\cdot)$, $\mathbf{x}$, and $\boldsymbol{\delta}$ respectively replaced by $\mathbf{w}_k$, $\mathbf{x}_{k+1}$, $\mathcal{H}_{\lambda/\mu}(g(\cdot))$, $\mathbf{x}_k$ and $\boldsymbol{\delta}_k$, we have:

$$\begin{aligned}
Q_\mu(\mathbf{x}_{k+1},\mathbf{w}_k) &\le Q_\mu(\mathcal{H}_{\lambda/\mu}(g(\mathbf{x}_k)),\mathbf{x}_k) \\
&\quad - 2\eta_k(1-\eta_k)\frac{[\boldsymbol{\delta}_k^T\mathbf{B}_\mu(\mathcal{H}_{\lambda/\mu}(g(\mathbf{x}_k))-\mathbf{x}_k)]^2}{\boldsymbol{\delta}_k\mathbf{B}_\mu\boldsymbol{\delta}_k} \\
&\le F(\mathbf{x}_k) - \sigma_k\alpha_k^2\boldsymbol{\delta}_k^T\mathbf{B}_\mu\boldsymbol{\delta}_k, \qquad (24)
\end{aligned}$$

where $\sigma_k = (1-\eta_k)/2\eta_k > 0$. The first term in (24) follows from the fact that:

$$Q_\mu(\mathcal{H}_{\lambda/\mu}(g(\mathbf{x}_k)),\mathbf{x}_k) = Q_\mu(\mathcal{P}_\mu(\mathbf{x}_k),\mathbf{x}_k) \le F(\mathbf{x}_k).$$

Now, noting that:

$$\begin{aligned}
Q_\mu(\mathbf{x}_{k+1},\mathbf{w}_k) &= F(\mathbf{x}_{k+1}) \\
&\quad + \frac{1}{2}(\mathbf{x}_{k+1}-\mathbf{w}_k)^T\mathbf{B}_\mu(\mathbf{x}_{k+1}-\mathbf{w}_k), \qquad (25)
\end{aligned}$$

which easily follows from basic linear algebra, (24) and (25) together imply that:

$$\begin{aligned}
F(\mathbf{x}_k) - F(\mathbf{x}_{k+1}) &\ge \sigma_k\alpha_k^2\boldsymbol{\delta}_k^T\mathbf{B}_\mu\boldsymbol{\delta}_k \\
&\quad + \frac{1}{2}(\mathbf{x}_{k+1}-\mathbf{w}_k)^T\mathbf{B}_\mu(\mathbf{x}_{k+1}-\mathbf{w}_k) \\
&\ge \rho\sigma_k\alpha_k^2\|\boldsymbol{\delta}_k\|_2^2 + \frac{\rho}{2}\|\mathbf{x}_{k+1}-\mathbf{w}_k\|_2^2, \qquad (26)
\end{aligned}$$

where $\rho > 0$ is the smallest eigenvalue of $\mathbf{B}_\mu \succ 0$. So, both terms on the right hand side in (26) are ≥ 0 for all k. As a result, due to (23) we can use the pinching/squeeze argument on (26) to establish that $\mathbf{x}_{k+1} - \mathbf{w}_k = \boldsymbol{\delta}_{k+1} - \alpha_k\boldsymbol{\delta}_k \to 0$ and $\alpha_k\boldsymbol{\delta}_k \to 0$ as $k\to\infty$. Consequently, $\boldsymbol{\delta}_k \to 0$ as $k\to\infty$, which completes the proof. □

Proof of Lemma 2: Firstly, from Lemma 1 we have that $\mathbf{x}_{k_n} - \mathbf{x}_{k_n-1} \to 0$ as $n\to\infty$. In the last paragraph in the proof of that lemma (above) we also have that $\alpha_k\boldsymbol{\delta}_k \to 0$, and so, $\alpha_{k_n}\boldsymbol{\delta}_{k_n} \to 0$ as $n\to\infty$. As a result, by the definition of $\mathbf{w}_{k_n}$, $\mathbf{w}_{k_n} \to \mathbf{w}_\bullet = \mathbf{x}_\bullet$. Then, by the continuity of $\nabla f(\cdot)$ we have that:

$$\mathbf{w}_{k_n} - \frac{1}{\mu}\nabla f(\mathbf{w}_{k_n}) \to \mathbf{x}_\bullet - \frac{1}{\mu}\nabla f(\mathbf{x}_\bullet).$$

Now, we need to show that, if $\mathbf{u}_{k_n} \to \mathbf{u}_\bullet$ as $n\to\infty$, then:

$$\mathcal{H}_{\lambda/\mu}(\mathbf{u}_{k_n}) \to \mathcal{H}_{\lambda/\mu}(\mathbf{u}_\bullet). \qquad (27)$$

Consider an arbitrary component of $\mathbf{u}_{k_n}$, say, $\mathbf{u}_{k_n}[i]$, in which case we must have $\mathbf{u}_{k_n}[i] \to \mathbf{u}_\bullet[i]$. Without loss of generality, assume $\mathbf{u}_\bullet[i] > 0$. Then, by the definition of $\mathcal{H}_{\lambda/\mu}(\cdot)$, there are two scenarios to consider:

$$\text{(a)}\ \ \mathbf{u}_\bullet[i] \neq \sqrt{2\lambda/\mu} \qquad \text{(b)}\ \ \mathbf{u}_\bullet[i] = \sqrt{2\lambda/\mu}.$$

Regarding (a): For a large enough $n = N$, we must either have $\mathbf{u}_{k_n}[i] < \sqrt{2\lambda/\mu}$ or $\mathbf{u}_{k_n}[i] > \sqrt{2\lambda/\mu}$ for all $n > N$, which implies:

$$\mathcal{H}_{\lambda/\mu}(\mathbf{u}_{k_n})[i] = \begin{cases} 0 & \text{if } \mathbf{u}_{k_n}[i] < \sqrt{2\lambda/\mu} \\ \mathbf{u}_{k_n}[i] & \text{if } \mathbf{u}_{k_n}[i] > \sqrt{2\lambda/\mu}, \end{cases} \qquad (28)$$

for all $n > N$. In (28), $\mathcal{H}_{\lambda/\mu}(\cdot)$ is a continuous function of $\mathbf{u}_{k_n}[i]$ in both cases, which immediately implies (27).

Regarding (b): In general, $\mathbf{u}_\bullet[i]$ could be reached in an oscillating fashion, i.e., for some n we can have $\mathbf{u}_{k_n}[i] < \sqrt{2\lambda/\mu}$ and for others $\mathbf{u}_{k_n}[i] > \sqrt{2\lambda/\mu}$. If this was the case for all $n \to \infty$ then $\mathcal{H}_{\lambda/\mu}(\mathbf{u}_{k_n})[i]$ would approach a limit set of two points $\{0, \sqrt{2\lambda/\mu}\}$. However, having $\mathbf{x}_{k_n} \to \mathbf{x}_\bullet$ and $\mathbf{x}_{k+1} - \mathbf{x}_k \to 0$ implies $\mathbf{x}_{k_n+1} \to \mathbf{x}_\bullet$. So, using the fact that:

$$\mathbf{x}_{k_n+1}[i] = \mathcal{H}_{\lambda/\mu}(\mathbf{u}_{k_n})[i] \tag{29}$$

$\mathcal{H}_{\lambda/\mu}(\mathbf{u}_{k_n})[i]$ must approach either 0 or $\sqrt{2\lambda/\mu}$. In other words, there has to exist a large enough $n = N$ such that $\mathbf{u}_\bullet[i]$ is approached either only from the left or the right for all $n > N$, i.e.,

(b_1) if $\mathbf{u}_{k_n}[i] < \mathbf{u}_\bullet[i] = \sqrt{2\lambda/\mu}$ for all $n > N$, from (28) we have $\mathcal{H}_{\lambda/\mu}(\mathbf{u}_{k_n})[i] \to 0$. So, noting that:

$$\mathcal{H}_{\lambda/\mu}(\mathbf{u}_{k_n})[i] - \mathbf{x}_{k_n} \overset{(29)}{=} \mathbf{x}_{k_n+1} - \mathbf{x}_{k_n} \to 0 \tag{30}$$

implies $\mathbf{x}_{k_n}[i] \to 0$. As a result, $\mathbf{x}_\bullet[i] = 0$, and using the definition of $\mathcal{H}_{\lambda/\mu}(\cdot)$, we have:

$$\begin{aligned}\mathcal{H}_{\lambda/\mu}(\mathbf{u}_\bullet)[i] &= \sqrt{2\lambda/\mu}\ \mathbb{I}(\mathbf{w}_\bullet[i] \neq 0)\\ &= \sqrt{2\lambda/\mu}\ \mathbb{I}(\mathbf{x}_\bullet[i] \neq 0) = \sqrt{2\lambda/\mu} \cdot 0\\ &= 0.\end{aligned}$$

Hence, (27) is satisfied.

(b_2) if $\mathbf{u}_{k_n}[i] > \mathbf{u}_\bullet[i] = \sqrt{2\lambda/\mu}$ for all $n > N$, from (28) we have $\mathcal{H}_{\lambda/\mu}(\mathbf{u}_{k_n})[i] \to \sqrt{2\lambda/\mu}$. So, (30) implies $\mathbf{x}_{k_n} \to \sqrt{2\lambda/\mu}$, and using the definition of $\mathcal{H}_{\lambda/\mu}(\cdot)$, we have:

$$\begin{aligned}\mathcal{H}_{\lambda/\mu}(\mathbf{u}_\bullet)[i] &= \sqrt{2\lambda/\mu}\ \mathbb{I}(\mathbf{w}_\bullet[i] \neq 0)\\ &= \sqrt{2\lambda/\mu}\ \mathbb{I}(\mathbf{x}_\bullet[i] \neq 0) = \sqrt{2\lambda/\mu} \cdot 1\\ &= \sqrt{2\lambda/\mu}.\end{aligned}$$

Hence, (27) is again satisfied.

Since i is arbitrary, the proof is complete. □

Proof of Lemma 3: The fixed points are obviously obtained by setting $\mathbf{x}_{k+1} = \mathbf{x}_k = \mathbf{x}_{k-1} = \mathbf{x}_\bullet$. So, any fixed point $\mathbf{x}_\bullet$ satisfies the equation:

$$\mathbf{x}_\bullet = \mathcal{H}_{\lambda/\mu}\left(\mathbf{x}_\bullet - \frac{1}{\mu}\nabla f(\mathbf{x}_\bullet)\right). \tag{31}$$

The result is established by using the definition of $\mathcal{H}_{\lambda/\mu}(g(\cdot))$ in Theorem 2. Namely, if $i \in \mathcal{Z}$ we easily obtain:

$$|(1/\mu)\nabla f(\mathbf{x}_\bullet)[i]| \leq \sqrt{2\lambda/\mu},$$

which reduces to C_1. If $i \in \mathcal{Z}^c$, we easily obtain:

$$\mathbf{x}_\bullet[i] = \mathbf{x}_\bullet[i] - (1/\mu)\nabla f(\mathbf{x}_\bullet)[i], \tag{32}$$

which reduces to C_2. However, since:

$$|\mathbf{x}_\bullet[i] - (1/\mu)\nabla f(\mathbf{x}_\bullet)[i]| \geq \sqrt{2\lambda/\mu}, \tag{33}$$

(32) and (33) together imply $|\mathbf{x}_\bullet[i]| \geq \sqrt{2\lambda/\mu}$, giving C_3. □

Proof of Lemma 4: Letting $\mathcal{Z} = \{i : \mathbf{x}_\bullet[i] = 0\}$ and $\mathcal{Z}^c = \{i : \mathbf{x}_\bullet[i] \neq 0\}$, it can easily be shown that $F(\mathbf{x}_\bullet + \mathbf{d}) = F(\mathbf{x}_\bullet) + \phi(\mathbf{d})$, where:

$$\begin{aligned}\phi(\mathbf{d}) &= \frac{1}{2}\|\mathbf{A}\mathbf{d}\|_2^2 + \mathbf{d}^T\nabla f(\mathbf{x}_\bullet) + \lambda\|\mathbf{x}_\bullet + \mathbf{d}\|_0 - \lambda\|\mathbf{x}_\bullet\|_0\\ &\geq \sum_{i\in\mathcal{Z}} \underbrace{\mathbf{d}[i]\nabla f(\mathbf{x}_\bullet)[i] + \lambda\mathbb{I}(\mathbf{d}[i] \neq 0)}_{=\phi_{\mathcal{Z}}(\mathbf{d}[i])}\\ &\quad + \sum_{i\in\mathcal{Z}^c} \underbrace{\mathbf{d}[i]\nabla f(\mathbf{x}_\bullet)[i] + \lambda\mathbb{I}(\mathbf{x}_\bullet[i] + \mathbf{d}[i] \neq 0) - \lambda}_{=\phi_{\mathcal{Z}^c}(\mathbf{d}[i])}.\end{aligned}$$

Now, $\phi_{\mathcal{Z}}(0) = 0$, so suppose $|\mathbf{d}[i]| \in \left(0, \lambda/\sqrt{2\lambda\mu}\right)$, $i \in \mathcal{Z}$. Then:

$$\phi_{\mathcal{Z}}(\mathbf{d}[i]) \geq -|\mathbf{d}[i]||\nabla f(\mathbf{x}_\bullet)[i]| + \lambda \geq -|\mathbf{d}[i]|\sqrt{2\lambda\mu} + \lambda > 0.$$

The second inequality in the above is due to C_1 in Lemma 3.

Lastly, note that $\phi_{\mathcal{Z}^c}(0) = 0$, and suppose $i \in \mathcal{Z}^c$. From C_2 in Lemma 3 we have $\nabla f(\mathbf{x}_\bullet)[i] = 0$. Thus, supposing $|\mathbf{d}[i]| \in (0, \sqrt{2\lambda/\mu})$, from C_2 we have $|\mathbf{d}[i]| < |\mathbf{x}_\bullet[i]|$ for all $i \in \mathcal{Z}^c$. Thus:

$$\mathbb{I}(\mathbf{x}_\bullet[i] + \mathbf{d}[i] \neq 0) = \mathbb{I}(\mathbf{x}_\bullet[i] \neq 0) = 1,$$

and so, $\phi_{\mathcal{Z}^c}(\mathbf{d}[i]) = 0$. Since $\frac{\lambda}{\sqrt{2\lambda\mu}} < \sqrt{\frac{2\lambda}{\mu}}$, the proof is complete after letting $\epsilon = \frac{\lambda}{\sqrt{2\lambda\mu}}$. □

Proof of Lemma 5: Since it is assumed that $\{\mathbf{x}_k\}_{k\geq 0}$ is bounded, the sequence $\{(\mathbf{x}_k, \mathbf{x}_{k+1})\}_{k\geq 0}$ is also bounded, and thus, has at least one limit point. Denoting one of these by $(\mathbf{x}_\bullet, \mathbf{x}_{\bullet\bullet})$, there exists a subsequence $\{(\mathbf{x}_{k_n}, \mathbf{x}_{k_n+1})\}_{n\geq 0}$ such that $(\mathbf{x}_{k_n}, \mathbf{x}_{k_n+1}) \to (\mathbf{x}_\bullet, \mathbf{x}_{\bullet\bullet})$ as $n \to \infty$. However, by Lemma 1 we must have $\mathbf{x}_{k_n} - \mathbf{x}_{k_n+1} \to 0$, which implies $\mathbf{x}_\bullet = \mathbf{x}_{\bullet\bullet}$. Consequently:

$$\mathbf{x}_{k_n+1} = \mathcal{H}_{\lambda/\mu}\left(\mathbf{w}_{k_n} - \frac{1}{\mu}\nabla f(\mathbf{w}_{k_n})\right) \to \mathbf{x}_\bullet, \tag{34}$$

recalling that $\mathbf{w}_{k_n} = \mathbf{x}_{k_n} + \alpha_{k_n}\boldsymbol{\delta}_{k_n}$ and $\boldsymbol{\delta}_{k_n} = \mathbf{x}_{k_n} - \mathbf{x}_{k_n-1}$. Furthermore, the convergence:

$$\mathcal{H}_{\lambda/\mu}\left(\mathbf{w}_{k_n} - \frac{1}{\mu}\nabla f(\mathbf{w}_{k_n})\right) \to \mathcal{H}_{\lambda/\mu}\left(\mathbf{x}_\bullet - \frac{1}{\mu}\nabla f(\mathbf{x}_\bullet)\right), \tag{35}$$

follows from Lemma 2. Equating the limits in (34) and (35) assures that $\mathbf{x}_\bullet$ satisfies the fixed point equation (31), making it a fixed point of the algorithm. This completes the proof. □

Proof of Theorem 3: By Lemma 1 and Ostrowski's result [23, Theorem 26.1], the bounded $\{\mathbf{x}_k\}_{k\geq 0}$ converges to a closed and connected set, i.e., the set of limit points form a closed and connected set. But, by Lemma 5 these limit points are fixed points, which by Lemma 4 are strict local minimizers. So, since the local minimizers form a discrete set the connected set of limit points can only contain one point, and so, the entire $\{\mathbf{x}_k\}_{k\geq 0}$ must converge to a single local minimizer. □

REFERENCES

[1] T. Blumensath, M. Yaghoobi, and M. E. Davies, "Iterative hard thresholding and l_0 regularisation," *IEEE ICASSP*, vol. 3, pp. 877–880, 2007.

[2] T. Blumensath and M. Davies, "Iterative thresholding for sparse approximations," *J. Fourier Anal. Appl.*, vol. 14, no. 5, pp. 629–654, 2008.

[3] R. Mazumder, J. Friedman, and T. Hastie, "SparseNet: Coordinate descent with non-convex penalties," *J. Am. Stat. Assoc.*, vol. 106, no. 495, pp. 1–38, 2011.

[4] K. Bredies, D. A. Lorenz, and S. Reiterer, "Minimization of non-smooth, non-convex functionals by iterative thresholding," *Tech. Rept.*, 2009, http://www.uni-graz.at/~bredies/publications.html.

[5] P. Tseng, "Convergence of block coordinate descent method for nondifferentiable minimization," *J. Optimiz. Theory App.*, vol. 109, no. 3, pp. 474–494, 2001.

[6] M. Nikolova, "Description of the minimisers of least squares regularized with l_0 norm. uniqueness of the global minimizer," *SIAM J. Imaging. Sci.*, vol. 6, no. 2, pp. 904–937, 2013.

[7] G. Marjanovic and V. Solo, "On exact l_q denoising," *IEEE ICASSP*, pp. 6068–6072, 2013.

[8] ——, "l_q sparsity penalized linear regression with cyclic descent," *IEEE T. Signal Proces.*, vol. 62, no. 6, pp. 1464–1475, 2014.

[9] A. Beck and M. Teboulle, "A fast iterative shrinkage-thresholding algorithm for linear inverse problems," *SIAM J. Imaging Sciences*, vol. 2, no. 1, pp. 183–202, 2009.

[10] S. J. Wright, R. D. Nowak, and M. Figueiredo, "Sparse reconstruction by separable approximation," *IEEE T. Signal Proces.*, vol. 57, no. 7, pp. 2479–2493, 2009.

[11] A. Beck and M. Teboulle, "Fast gradient-based algorithms for constrained total variation image denoising and deblurring," *IEEE T. Image Process.*, vol. 18, no. 11, pp. 2419–2134, 2009.

[12] M. Figueiredo, J. M. Bioucas-Dias, and R. D. Nowak, "Majorization-minimization algorithms for wavelet-based image restoration," *IEEE T. Image Process.*, vol. 16, no. 12, pp. 2980–2991, 2007.

[13] Y. Nesterov, "A method of solving a convex programming problem with convergence rate $\mathcal{O}(1/k^2)$," *Soviet Math. Doklady*, vol. 27, pp. 372–376, 1983.

[14] U. J. Schwarz, "Mathematical-statistical description of the iterative beam removing technique (method clean)," *Astron. Astrophys.*, vol. 65, pp. 345–356, 1978.

[15] H. Akaike, "Information theory and an extension of the maximum likelihood principle," *2nd International Symposium on Information Theory*, pp. 267–281, 1973.

[16] M. Stone, "Cross-validatory choice and assessment of statistical predictions (with Discussion)," *J. R. Statist. Soc. B*, vol. 39, pp. 111–147, 1974.

[17] P. Craven and G. Wahba, "Smoothing noisy data with spline functions: Estimating the correct degree of smoothing by the method of generalized cross-validation," *Numer. Math.*, vol. 31, pp. 377–403, 1979.

[18] J. Chen and Z. Chen, "Extended Bayesian information criteria for model selection with large model spaces," *Biometrika*, vol. 95, no. 3, pp. 759–771, 2008.

[19] K. W. Broman and T. P. Speed, "A model selection approach for the identification of quantitative trait loci in experimental crosses," *J. R. Statist. Soc. B*, vol. 64, pp. 641–656, 2002.

[20] D. Siegmund, "Model selection in irregular problems: Application to mapping quantitative trait loci," *Biometrika*, vol. 91, pp. 785–800, 2004.

[21] M. Bogdan, R. Doerge, and J. K. Ghosh, "Modifying the Schwarz Bayesian information criterion to locate multiple interacting quantitative trait loci," *Genetics*, vol. 167, pp. 989–999, 2004.

[22] G. Marjanovic and A. O. Hero, "On l_q estimation of sparse inverse covariance," *IEEE ICASSP*, 2014.

[23] A. M. Ostrowski, *Solutions of Equations in Euclidean and Banach Spaces*. New York: Academic Press, 1973.